

A Conversational Approach to Process-oriented Case-based Reasoning

Christian Zeyen, Gilbert Müller and Ralph Bergmann

University of Trier, Business Information Systems II

zeyen@uni-trier.de, muellerg@uni-trier.de, bergmann@uni-trier.de

Abstract

Process-oriented case-based reasoning (POCBR) supports workflow modeling by retrieving and adapting workflows that have proved useful in the past. Current approaches typically require users to specify detailed queries, which can be a demanding task. Conversational case-based reasoning (CCBR) particularly addresses this problem by proposing methods that incrementally elicit the relevant features of the target problem in an interactive dialog. However, no CCBR approaches exist that are applicable for workflow cases that go beyond attribute-value representations such as labeled graphs. This paper closes this gap and presents a conversational POCBR approach (C-POCBR) in which questions related to structural properties of the workflow cases are generated automatically. An evaluation with cooking workflows indicates that C-POCBR can reduce the communication effort for users during retrieval.

1 Introduction and Foundations

Process-oriented case-based reasoning (POCBR) [Minor *et al.*, 2014] addresses the integration of case-based reasoning (CBR) [Richter and Weber, 2013; Bergmann, 2002] with process-aware information systems [Van Der Aalst, 2013] to provide experience-based support for various tasks such as workflow modeling, monitoring, analysis, or execution. In this paper, we focus on workflow modeling by reuse of best-practice workflows (workflow cases) from a repository (case base). Thus, we aim at retrieving already available workflows that can be adapted for new purposes and circumstances. A case in POCBR is usually a workflow or process description expressing procedural experiential knowledge. Broadly speaking, a workflow describes a logical or chronological order of tasks that exchange physical products or data to reach a certain outcome – the workflow output [Van Der Aalst, 2013]. Cooking recipes can be considered as a simple form of process descriptions. In cooking workflows (see Figure 1), tasks represent required cooking steps and exchange ingredients in order to produce a certain dish. We describe workflows as semantically labeled directed graphs by adopting the representation by Bergmann and Gil 2014. The semantic labels of




Figure 1: Example of a Cooking Workflow

nodes are structured hierarchically in a taxonomy of ingredients and a taxonomy of cooking steps. In order to retrieve workflows in traditional POCBR, a fully elaborated query describing the users' requirements is typically assumed to be available from the beginning [Bergmann and Gil, 2014]. For example, in the process-oriented query language (POQL) [Müller and Bergmann, 2015c] a query consists of a query workflow and several restriction workflows. A query workflow represents properties the searched workflow should fulfill while restriction workflows represent undesired situations that should be avoided. However, in practice, users may only have a vague idea of the workflow they desire or they lack detailed domain knowledge and thus have serious difficulties to provide a precise query. In order to obtain a reusable workflow, similarity search or process model querying can be applied [Dijkman *et al.*, 2012]. Outside of POCBR, various query languages have been proposed [Awad, 2007; Beeri *et al.*, 2006; Störrle and Acretoai, 2013], which are used in visual query editors to formulate graph-based queries. Typically, matching workflows from a repository are obtained by applying graph edit measures [Dijkman *et al.*, 2009] or graph/subgraph similarity measures [Bergmann and Gil, 2014].

While in many CBR applications a complete description of the target problem (i.e. the query) is assumed to be available in advance, conversational CBR (CCBR) [Aha *et al.*, 2001; Mcsherry, 2003; Aha *et al.*, 2005] particularly addresses the interactive nature of problem solving. CCBR approaches include methods which incrementally elicit the relevant features of the target problem in an interactive dialog, often with the aim of minimizing the communication effort for the user. The basic assumption behind CCBR is that guided question answering requires less domain expertise than providing de-

tailed queries from scratch. A goal in a conversation is to pose relevant questions, potentially suitable to elaborate the query automatically and to retrieve the most useful case efficiently. To perform the dialog, the case representation in CCBR is enriched with an additional set of question-answer pairs stated in natural language. Thus, case authoring can become more demanding, since suitable questions need to be formulated. Hence, the automatic creation of questions is desirable and often achieved by deriving questions from case attributes. CCBR research focuses on enhancing case representation to include knowledge relevant for the questioning strategy [Gu and Aamodt, 2005], methods for question selection, and methods for dialog inferencing and termination [Kohlmaier *et al.*, 2001; Mcsherry, 2003; Aha *et al.*, 2005; Gu and Aamodt, 2005]. Our research is based on the similarity variance measure proposed by Kohlmaier *et al.* 2001 which prefers questions, whose answers most probably have the highest influence on the similarity distribution of the most similar cases. CCBR finds its application mostly in analytical applications such as sequential diagnosis, customer helpdesk support, or product recommendation. Only very few approaches have been proposed that address synthetic applications [Leake and Wilson, 1999; Muñoz-Avila *et al.*, 1999; Weber *et al.*, 2004]. Today, no CCBR approaches exist so far that automatically elicit questions from case descriptions that go beyond attribute-value representation to construct queries for retrieval. In particular, no such approach is applicable for workflow representations as required for POCBR.

This paper presents a new conversational POCBR (C-POCBR) approach that deals with workflows represented as labeled graphs. It is an abridged version of our full paper [Zeyen *et al.*, 2017]. The C-POCBR retrieval considers the structural properties of workflows. Questions related to those properties are automatically constructed based on extracted workflow fragments and a respective question selection strategy. With this approach, we aim to reduce the effort and the required expertise for the definition of queries in POCBR. We illustrate and evaluate the approach with cooking workflows.

In the following, section 2 describes our C-POCBR approach in more detail. An experimental evaluation is presented in section 3 while section 4 summarizes our findings and discusses future work.

2 A Conversational POCBR Approach

Based on the idea of CCBR, we now present a new approach that is particularly tailored to POCBR and thus named conversational POCBR (C-POCBR). In a nutshell, users are guided through the query process by a sequence of questions about their desired workflows. The more questions are answered, the more knowledge about desired and undesired properties is available, which is stored in an internal query for retrieval. A major focus is put on the automatic creation of questions to avoid that they need to be specified manually. For this purpose, we consider workflow fragments as characteristic properties of a workflow, which we refer to as features. The basic idea is to automatically extract features from the workflows stored in the case base, which are then used as the subject of questions. In order to conduct efficient conversations, we

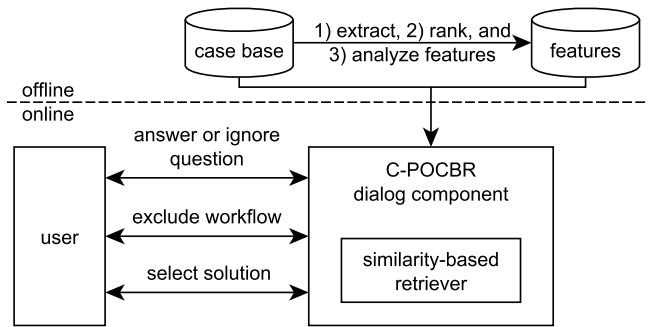


Figure 2: Conversational POCBR Process

rank features by their ability to distinguish workflows from one another. Furthermore, identified relations between features enable to generate coherent follow-up questions and to infer irrelevant features based on already answered questions.

The conversational POCBR process is divided into two phases (see Figure 2). The *offline phase* comprises pre-computations for the initial setup. During this phase, extraction, ranking, and analysis of features takes place. Subsequently, the actual conversation is conducted in the *online phase*.

2.1 Offline Phase

At first, features are extracted based on the graph representation of the workflows. As those features will occur in the questions posed to the user, they must be as simple and understandable as possible. For this purpose, we consider various design guidelines investigated in related work [Aha *et al.*, 2001; Kohlmaier *et al.*, 2001; Richter and Weber, 2013]. In principle, a feature can be any fragment of a workflow. The smallest possible feature consists of a single workflow item. This can be a single node such as a data or a task node. More complex features can be created by extracting partial workflows. To derive questions on a more general level of detail, we apply a generalization algorithm [Müller and Bergmann, 2015b], which generalizes semantic labels based on the domain taxonomies. We extract two different kinds of features from each workflow in the case base: *specific* and *generalized feature nodes* as well as *specific* and *generalized feature workflows*, i.e., partial workflows that consist of at least one task and one data node.

In a second step, features are sorted in descending order by their ability to distinguish workflows from one another. We adopt the *simVar* measure by Kohlmaier *et al.* 2001, which utilizes the similarity variance as a ranking criterion. It estimates the variance of the similarity of the most similar cases assuming that the value of the respective feature in the query is known. Consequently, we prefer features with a higher *simVar* value. Due to the POQL query, the user can either select a feature as desired or undesired during the conversation. Therefore, the similarity variance is pre-computed for both situations. To calculate *simVar*, all similarities between the extracted features and the workflows stored in the case base must be computed. Each feature is added into the query part and the restriction part of an empty query, respectively.

Then, for both queries, the similarities to each workflow from the case base are computed and cached. Two similarity measures are required, one measure to assess the similarity of a case w.r.t. the query workflow and one measure to assess the similarity w.r.t. the restriction workflows. For the query workflow we use a similarity measure by Bergmann and Gil 2014, which treats the similarity computation as an optimization problem solved by searching the best possible mapping of all nodes and edges of the query workflow to those of each workflow case. The core of the similarity model is a local similarity measure for semantic descriptions. In our example domain, similarity values between semantic labels are derived from the data and task taxonomy that reflect the closeness of the concepts. For the restriction workflows in a query we use a binary measure that evaluates whether a workflow case violates a given restriction or not. Both measures are combined to produce the overall ranking of the cases during retrieval.

In the next step, relations between features are analyzed. For each feature all related features are determined. Related features share a common partial workflow (including generalizations). We differentiate features by their number of nodes and by their generality of nodes. For example, for the feature workflow $f_1 = \{slice, ham\}$, the related feature $g_1 = \{cut, meat\}$ is *more general* and *equally large* while the feature $g_2 = \{parma-ham\}$ is *more specific* and *smaller*.

2.2 Online Phase

In the online phase, the C-POCBBR dialog component iteratively creates and displays questions until a desired workflow is found or until stopping criteria are fulfilled. The dialog starts with an empty query¹. Initially, the set of candidate features CF , i.e., relevant features to be asked in a question, is the full set of extracted features. The initial set of candidate workflows CW , i.e., the solution space, encompasses the whole case base.

In the main loop, the dialog component selects a question based on the candidate features CF considering the $simVar$ scoring, the previously answered questions, as well as the feature relations. Each question involves one or in certain cases several candidate features and gives the user four options to react:

1. *Ignoring a question:* The feature that is the subject of the question as well as larger or more specific related features are removed from the set of candidate features and the next best question is displayed.
2. *Answering a question:* The query is extended by either a desired or an undesired feature depending on the given answer and a similarity-based retrieval with the extended query on the current set of candidate workflows CW is performed. After each retrieval, workflows that are less similar than the average of all workflow's similarities are removed from CW and thus are not included in subsequent retrievals. By this means, only the most suitable workflows are retained with respect to the current query. The workflow with the highest similarity from CW is

¹In principle an initial pre-modeled query could be used as well, but we have not yet investigated this option.

displayed to the user. In addition, only features contained in the candidate workflows CW remain in the set of candidate features CF and their ranking is updated according to $simVar$ by using CW instead of CB . By this means, it is ensured that only relevant questions are posed and with each question answered, CW and CF are further reduced.

3. *Excluding a suggested workflow:* This removes the workflow from the case base CB (only temporary for this dialog) and triggers a new retrieval on the whole case base. As a consequence, it is likely that more candidate workflows CW than before exist because of the lower average similarity of all workflows to the current query. Consequently, more candidate features CF may become available.
4. *Selecting a solution:* With a desired workflow selected by the user the retrieval terminates successfully.

We now describe in more detail the applied question selection method. The question sequence mainly consists of three types of questions. At the beginning of a conversation the highest ranked feature from the candidate features is suggested in a feature question (FQ). This type of question is not related to previously suggested features and it will be asked as long as the user rejects a suggestion. If the user selects a feature as desired in a FQ , a first follow-up question, i.e., a specialization question (SQ), is posed suggesting one or (if available) several equally large but more specific features. Again, the features are sorted by their $simVar$ value. This type of question is repeated as long as the user chooses specializations. Following the SQs , an enlargement question (EQ) is displayed to the user that suggests larger and not more general features than the previously selected and/or specialized feature. If no more EQs are available, the next initial FQ is selected, addressing a new and potentially unrelated subject matter.

When the set of candidate features CF is updated due to an ignored or answered question, irrelevant features can be inferred based on the relations between features. If a question is marked as irrelevant, all the related features (e.g., more specific and larger features) are marked as irrelevant, too. If suggested features are selected as undesired, they are added to the restriction part of the current query and related irrelevant features are no longer considered as candidate features, to prevent the system from repetitively asking the user what she does not like. If a feature is marked as desired, also related features such as more general features are removed from CF . If a user chooses a specialization or an enlargement, the target feature that is already present in the query is replaced with the new feature. In this event, related features of the target feature without those that are still relevant for the new feature are removed from CF .

3 Evaluation

We now describe the evaluation comparing the presented C-POCBBR approach with a traditional POCBBR approach in which the user models a POQL query manually using a query editor. The evaluation aims at testing three hypotheses:

- H1** The desired workflow is retrieved with C-POCBR when all questions are answered correctly.
- H2** The C-POCBR dialog enables users to retrieve the desired workflow.
- H3** C-POCBR reduces the communication effort required to retrieve the desired workflow.

3.1 Evaluation Setup

For the experiments, we use the already existing CookingCAKE system [Müller and Bergmann, 2015a], which is part of the CAKE framework². It already includes a graphical POQL editor, which is used as implementation of the POCBR approach. In addition, we implemented the C-POCBR approach³ as an extension of CookingCAKE. We use a case base of 61 cooking workflows from which we derived 60 search scenarios in a semi-automatic process. Each search scenario unambiguously describes a target workflow from the case base in plain text and is given to the users.

3.2 Experimental Evaluation

Hypothesis H1 is tested using an experiment with a simulated user, which automatically answers the posed questions of the C-POCBR approach correctly. We adopt the methodology by Aha et al. [Aha et al., 2001], who evaluate a conversational retrieval with a leave-one-in cross validation. Consequently, in each search scenario the corresponding target workflow remains in the case base. It turned out that the target workflow is retrieved in each of the 60 search scenarios, which fully confirms hypothesis H1. In average, 10.25 questions were asked in the dialog. Hypotheses H2 and H3 are tested in experiments with eight human users who evaluated both approaches on the basis of four textual search scenarios of similar size.

Table 1 summarizes average results over all successful queries and all users for the POCBR and the C-POCBR approach. The *number of successful conversations* shows that only a few runs were not successful. Thus, hypothesis H2 can be confirmed. To assess the communication effort, the conversation time used in the POCBR and the C-POCBR approach were compared. In addition, the number of questions posed in the C-POCBR approach were determined. The *total conversation time* is the time span from the start of the conversation (in C-POCBR) or the start of the query input (in POCBR) until the desired workflow is retrieved and identified by the user. For POCBR and C-POCBR those time spans are quite comparable. We discovered that users following the POCBR approach tend to completely model the

	POCBR	C-POCBR
number of successful conversations	15/16	14/16
total conversation time	5:34 min.	5:40 min. (30 questions)
required conversation time	4:46 min.	2:16 min. (9 questions)

Table 1: Experimental Results: Avg. Values Across all Successful Retrievals and Users

²See cake.wi2.uni-trier.de

³See cookingcake.wi2.uni-trier.de/conversation

given query scenario, before they start the retrieval for the first time. In C-POCBR the users follow the dialog and investigate the presented workflow. When analyzing these results in more detail, we found out that quite often the dialog could have been terminated earlier as the user did not recognize the target workflow when it was displayed the first time. We analyzed this effect in detail and determined the *required conversation time* and the required number of questions the user was asked during this period. We can see that if users would have checked the displayed workflows more thoroughly, the C-POCBR approach could have been more than twice as fast as the POCBR approach. This issue shows that the workflow presentation in the C-POCBR implementation needs additional explanation functions that better allow the user to identify how the presented workflow relates to the answers of her query. However, with respect to the dialog component, we consider hypothesis H3 at least partially confirmed. Regarding the results obtained from questionnaires, the average ratings over 16 conversations assess the majority of the posed questions to be comprehensible and relevant. The results indicate that the automatic creation of questions provides useful questions for the conversation.

4 Conclusions and Future Work

We presented a novel approach to conversational POCBR that conducts an interactive dialog with users to facilitate the retrieval of workflows. To save effort for defining suitable questions, a method for the automatic creation of questions based on extracted features was described. Our work showed that those features are meaningful subjects of questions and that they are suitable to distinguish workflow cases from one another. The quality and performance of conversations was improved by ranking, analyzing, and selecting relevant features. We evaluated the approach with a simulated and real users and showed that the desired workflow can be found in a straight-forward manner. Furthermore, our results indicate that C-POCBR retrieval is able to reduce the communication effort for users.

In future work, it should be investigated how the presentation and explanation of workflows and features can be improved. It is desirable to evaluate the approach in domains with more complex workflows. In such domains, we assume that the conversation more strongly outperforms the traditional retrieval. In addition to POCBR domains, we assume the questioning strategy presented to be also applicable more broadly in CCB with complex case representations, provided that feature vocabularies are organized in a hierarchy and co-occurring features are identified. Also, future work could investigate how adaptability of workflows can be considered during a conversation. By this means, interactive retrieval could be combined with interactive adaptation to provide an even more powerful problem solver for users.

Acknowledgments

This work was funded by the German Research Foundation (DFG), project number BE 1373/3-3.

References

- [Aha *et al.*, 2001] David W. Aha, Leonard A. Breslow, and Héctor Muñoz-Avila. Conversational case-based reasoning. *Applied Intelligence*, 14(1):9–32, 2001.
- [Aha *et al.*, 2005] David W. Aha, David McSherry, and Qiang Yang. Advances in conversational case-based reasoning. *Knowledge Engineering Review*, 20(3):247–254, 2005.
- [Awad, 2007] Ahmed Awad. BPMN-Q: a language to query business processes. In *Enterprise Modelling and Information Systems Architectures - Concepts and Applications*, volume P-119 of *LNI*, pages 115–128. GI, 2007.
- [Beeri *et al.*, 2006] Catriel Beeri, Anat Eyal, Simon Kamenskovich, and Tova Milo. Querying business processes. In Umeshwar Dayal, Kyu-Young Whang, David B. Lomet, Gustavo Alonso, Guy M. Lohman, Martin L. Kersten, Sang Kyun Cha, and Young-Kuk Kim, editors, *Proceedings of the 32nd International Conference on Very Large Data Bases, 2006*, pages 343–354. ACM, 2006.
- [Bergmann and Gil, 2014] Ralph Bergmann and Yolanda Gil. Similarity assessment and efficient retrieval of semantic workflows. *Inf. Syst.*, 40:115–127, 2014.
- [Bergmann, 2002] Ralph Bergmann. *Experience Management: Foundations, Development Methodology, and Internet-based Applications*. Springer-Verlag, Berlin, Heidelberg, 2002.
- [Dijkman *et al.*, 2009] Remco M. Dijkman, Marlon Dumas, and Luciano García-Bañuelos. Graph matching algorithms for business process model similarity search. In Umeshwar Dayal, Johann Eder, Jana Koehler, and Hajo A. Reijers, editors, *Business Process Management, 7th International Conference, BPM 2009*, volume 5701 of *LNCS*, pages 48–63. Springer, 2009.
- [Dijkman *et al.*, 2012] Remco M. Dijkman, Marcello La Rosa, and Hajo A. Reijers. Managing large collections of business process models - current techniques and challenges. *Computers in Industry*, 63(2):91–97, 2012.
- [Gu and Aamodt, 2005] Mingyang Gu and Agnar Aamodt. A knowledge-intensive method for conversational CBR. In Héctor Muñoz-Ávila and Francesco Ricci, editors, *Case-Based Reasoning, Research and Development, ICCBR 2005*, volume 3620 of *LNCS*, pages 296–311. Springer, 2005.
- [Kohlmaier *et al.*, 2001] Andreas Kohlmaier, Sascha Schmitt, and Ralph Bergmann. A similarity-based approach to attribute selection in user-adaptive sales dialogs. In David W. Aha and Ian Watson, editors, *Case-Based Reasoning Research and Development, ICCBR 2001*, volume 2080 of *LNCS*, pages 306–320. Springer, 2001.
- [Leake and Wilson, 1999] David B. Leake and David C. Wilson. Combining CBR with interactive knowledge acquisition, manipulation and reuse. In *Case-Based Reasoning and Development, ICCBR '99*, pages 203–217. Springer, 1999.
- [Mcsherry, 2003] David Mcsherry. Increasing dialogue efficiency in case-based reasoning without loss of solution quality. In *IJCAI-03*, pages 121–126. Morgan Kaufmann, 2003.
- [Minor *et al.*, 2014] Mirjam Minor, Stefania Montani, and Juan A. Recio-Garcia. Process-oriented case-based reasoning. *Information Systems*, 40:103 – 105, 2014.
- [Müller and Bergmann, 2015a] Gilbert Müller and Ralph Bergmann. CookingCAKE: A framework for the adaptation of cooking recipes represented as workflows. In Joseph Kendall-Morwick, editor, *Workshop Proceedings from (ICCBR 2015)*, volume 1520 of *CEUR*, pages 221–232. CEUR-WS.org, 2015.
- [Müller and Bergmann, 2015b] Gilbert Müller and Ralph Bergmann. Generalization of workflows in process-oriented case-based reasoning. In Ingrid Russell and William Eberle, editors, *Proceedings of the 28th Int. Florida Artificial Intelligence Research Society Conference, FLAIRS 2015*, pages 391–396. AAAI Press, 2015.
- [Müller and Bergmann, 2015c] Gilbert Müller and Ralph Bergmann. POQL: A new query language for process-oriented case-based reasoning. In Ralph Bergmann, Sebastian Görg, and Gilbert Müller, editors, *Proceedings of the LWA 2015*, volume 1458 of *CEUR Workshop Proceedings*, pages 247–255. CEUR-WS.org, 2015.
- [Muñoz-Avila *et al.*, 1999] Héctor Muñoz-Avila, David W. Aha, Leonard A. Breslow, and Dana S. Nau. HICAP: an interactive case-based planning architecture and its application to noncombatant evacuation operations. In *Proc. of the 16th Nat. Conf. on Artif. Intell., AAAI/IAAI 1999*, pages 870–875. AAAI Press, 1999.
- [Richter and Weber, 2013] Michael M. Richter and Rosina O. Weber. *Case-Based Reasoning - A Textbook*. Springer, 2013.
- [Störrle and Acretoai, 2013] Harald Störrle and Vlad Acretoai. Querying business process models with VMQL. In *Proceedings of the 5th ACM SIGCHI Annual International Workshop on Behaviour Modelling - Foundations and Applications, BMFA '13*, pages 4:1–4:10, New York, 2013. ACM.
- [Van Der Aalst, 2013] Wil MP Van Der Aalst. Business process management: a comprehensive survey. *ISRN Software Engineering*, 2013, 2013.
- [Weber *et al.*, 2004] Barbara Weber, Werner Wild, and Ruth Breu. Cbrflow: Enabling adaptive workflow management through conversational case-based reasoning. In Peter Funk and Pedro A. González Calero, editors, *Advances in Case-Based Reasoning*, volume 3155 of *LNCS*, pages 434–448. Springer, 2004.
- [Zeyen *et al.*, 2017] Christian Zeyen, Gilbert Müller, and Ralph Bergmann. Conversational process-oriented case-based reasoning. In *Proceedings of ICCBR 2017*, *LNCS*, pages 403–419. Springer, 2017.